

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of)	
)	
Stefan JOHANSSON et al.)	Group Art Unit: 2171
)	
Application No.: 09/461,110)	Examiner: Unassigned
)	
Filed: December 14, 1999)	
)	
For: Method and Arrangement in a)	
Distributed System)	

CLAIM FOR CONVENTION PRIORITY

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

The benefit of the filing date of the following prior foreign application in the following foreign country is hereby requested, and the right of priority provided in 35 U.S.C. § 119 is hereby claimed:

Sweden Patent Application No. 9804333-4

Filed: December 14, 1998

In support of this claim, enclosed is a certified copy of the prior foreign application. The prior foreign application was referred to in the oath or declaration. Acknowledgment of receipt of the certified copy is requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

Date: April 25, 2002

By: Michael G. Savage
Michael G. Savage
Registration No. 32,596

P.O. Box 1404
Alexandria, Virginia 22313-1404
(919) 941-9240

PRV

PATENT- OCH REGISTRERINGSVERKET
Patentavdelningen

Intyg Certificate



Härmed intygas att bifogade kopior överensstämmer med de handlingar som ursprungligen ingivits till Patent- och registreringsverket i nedannämnda ansökan.

This is to certify that the annexed is a true copy of the documents as originally filed with the Patent- and Registration Office in connection with the following patent application.

(71) *Sökande* *Telefonaktiebolaget L M Ericsson, Stockholm SE*
Applicant (s)

(21) *Patentansökningsnummer* *9804333-4*
Patent application number

(86) *Ingivningsdatum* *1998-12-15*
Date of filing

Stockholm, 1999-10-27

För Patent- och registreringsverket
For the Patent- and Registration Office


Emma Johansson

Avgift
Fee *170:-*

Ink. t. Patent- och reg.verket

1998 -12- 1 5

METHOD AND ARRANGEMENT IN A DISTRIBUTED SYSTEM

TECHNICAL FIELD OF THE INVENTION

The invention relates to a method and an arrangement in a distributed system comprising a receiving unit and a source unit. More specifically the invention relates to a method
5 and arrangement for synchronizing configuring data at the receiving unit with corresponding source configuring data stored at the source unit.

DESCRIPTION OF RELATED ART

In many distributed systems there is a need of a mechanism
10 for synchronizing configuring data at a receiving unit with corresponding source configuring data stored at a source unit in order to ensure that the configuring data at the receiving unit is consistent with the corresponding source configuring data stored at a source unit.

15 One example of such a distributed system is the CMS8800 system, i.e. Ericsson's cellular system for D-AMPS, where a number of base stations are connected to a mobile services switching centre. In present CMS8800 systems configuring data at a base station is maintained consistent with
20 corresponding source configuring data at the mobile services switching centre by downloading a complete copy of the source configuring data that relates to the base station upon each system restart that takes place after a period of interruption of system operation. This approach ensures that
25 after a system restart the configuring data at the base station is consistent with the corresponding source configuring data at the mobile services switching centre. However, often when a system restart occurs, the configuring data at the base station is in fact already consistent with
30 the source configuring data at the mobile services switching centre. Since downloading configuring data to a number of base stations is time consuming and causes significant loads in primarily the mobile services switching centre and

communication links interconnecting the mobile services switching centre and the base stations, it would be preferable to have a method of maintaining said consistency which eliminates or at least lessens the number of occasions
5 when configuring data is downloaded from the mobile services switching centre when there is no inconsistency between the configuring data at the base station and the source configuring data at the mobile services switching centre.

The european patent application EP 665 670 discloses a
10 method and apparatus for decreasing the time required to update files located at a remote computer. The method identifies and isolates the difference between a reference file in a receiveing computer and a source file in a sending computer. Only said differences are transmitted to the
15 receiving computer.

EP 665 670 also discloses a method for synchronizing data at a receiving unit with data at a source unit. Multiple reference keys corresponding to groups of data stored at the receiving unit are determined and transmitted to the source
20 unit. A source key corresponding to a group of data in the source unit is determined and compared with the multiple reference keys. If the source key does not match any of the reference keys, data is transmitted from the source unit to the receiving unit. If the source key matches a reference
25 key, a control signal is transmitted to the receiving unit which causes the receiving unit to use data at the corresponding unit corresponding to the matched reference key. The steps of calculating a source key, comparing the source key to the multiple reference keys and transmitting
30 data/control signals are repeated for additional groups of source data in the source unit until the data at the receiving unit has been synchronized with the data at the source unit.

In the synchronization method disclosed in EP 665 670, the
35 source unit determines which data groups at the receiving

1998 -12- 15

3

unit that need to be updated. In a situation when the source unit is connected to a plurality of receiving units, the determination of which data groups that need to be updated in the receiving units causes a significant processing load at the source unit.

Another problem with the synchronization method disclosed in EP 665 670 is that the method requires the source unit to calculate source keys for the data groups in the source unit each time synchronization is performed. This results in further processing load at the source unit if synchronization is repeatedly performed and in particular when the source unit is connected to a plurality of receiving units.

SUMMARY OF THE INVENTION

The problem dealt with by the present invention is to provide an efficient way of synchronizing configuring data at a receiving unit with corresponding source configuring data.

The problem is solved essentially by a method and arrangement in which the receiving unit determines whether the content of the configuring data at the receiving unit matches corresponding reference checksums and requests downloading of data from the source unit upon detecting inconsistency.

More specifically, the problem is solved in the following manner. The configuring data and the source configuring data are each arranged in one or preferably several groups of data. Reference checksums are calculated for each data group. The receiving unit determines whether the content in each data group of the configuring data at the receiving unit matches the corresponding reference checksum. The receiving unit requests the source unit to transfer copies to the receiving unit of the source configuring data in

1998 -12- 15

4

those data groups for which a mismatch is detected. The requested data is downloaded from the source unit to the receiving unit.

5 A general object of the invention is to provide an efficient way of synchronizing configuring data at a receiving unit with corresponding source configuring data at a source unit.

10 A more specific object of the invention is to provide a way of synchronizing data in which only those data groups of the configuring data at the receiving unit which are suspected of being corrupted need to be downloaded anew from the source unit.

15 Another object of the invention is to provide a way of synchronizing data in which the source unit is relieved of the task of determining which data groups of the configuring data at the receiving unit that need to be updated due to the data groups being corrupted.

20 Still another object of the invention is to provide a way of synchronizing data in which the source unit is not required to calculate reference checksums anew each time synchronization is performed.

25 A general advantage afforded by the invention is that configuring data at a receiving unit can be synchronized with corresponding source configuring data at a source unit in an efficient way, in particular when synchronization is repeatedly performed and/or the source unit is connected to a plurality of receiving units.

30 A more specific advantage offered by the invention is that only those data groups of the configuring data at the receiving unit which are suspected of being corrupted need to be downloaded anew from the source unit upon synchronization.

Another advantage of the invention is that the source unit upon synchronization is relieved of the task of determining which data groups of the configuring data at the receiving unit that need to be updated due to the data groups being corrupted.

Still another advantage provided by the invention is that the source unit does not need to calculate reference checksums anew each time synchronization is performed.

The invention will now be described in more detail with reference to exemplifying embodiments thereof and also with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic view of a distributed system in which the invention is applied.

Fig. 2 is a flow chart illustrating a basic method according to the invention.

Fig. 3 is a view of a cellular network.

Fig. 4 is a block diagram illustrating a receiving unit in the form of a base station.

Fig. 5A-5B are diagrams illustrating data groups of configuring data at the base station of Fig. 4.

Fig. 6 is a functional block diagram illustrating function blocks related to the present invention in a mobile services switching centre and a base station connected to the mobile services switching centre.

Fig. 7 is a flow chart illustrating a preferred embodiment of the inventive method applied to the mobile services switching centre and the base station illustrated in Fig. 6.

DETAILED DESCRIPTION OF THE EMBODIMENTS

Fig. 3 illustrates a cellular network NET1 and a number of mobile stations MS1-MS3. The cellular network NET1 comprises a mobile services switching centre MSC1 and base stations BS1-BS5 connected to the mobile services switching centre MSC1. The base stations BS1-BS5 provide radio coverage in a geographical area served by the mobile services switching centre MSC1. The geographical area is divided into a number of cells C1-C5. In each cell C1-C5 radio coverage is provided by one of the base stations BS1-BS5 respectively. The mobile services switching centre MSC1 is responsible for switching calls to and from mobile stations MS1-MS3 located in the geographical area served by the mobile services switching centre MSC1. Note that the cellular network NET1 in Fig. 3 is an accurate but simplified illustration of a cellular network. A typical cellular network comprises several mobile services switching centres, a greater number of base stations as well as other types of nodes such as home location registers.

In the cellular network NET1, the mobile services switching centre MSC1 downloads configuring data to the base stations BS1-BS5 connected to the mobile services switching centre. The configuring data downloaded to each base station BS1-BS5 comprises different parameters being set to enable proper operation of each base station BS1-BS5. Examples of such parameters are those defining physical radio channels and logical channels allocated to a cell. A common characteristic of all the configuring data downloaded from the mobile services switching centre MSC1 is that it should be treated as read-only data at each base station BS1-BS5, i.e. the base station BS1-BS5 should not modify any part of the configuring data. Thus the configuring data at the base station BS1-BS5 should be consistent with corresponding source configuring data stored at the mobile services switching centre MSC1. However, due to disturbances of a base station BS1, the configuring data at the base station BS1 may become corrupted. Hence there is a need of a method

and arrangement for synchronizing the configuring data at the base stations BS1-BS5 with the corresponding source configuring data at the mobile services switching centre MSC1.

- 5 The present invention provides such a method and arrangement for data synchronization.

Fig. 1 is a schematic view of a distributed system 100 in which the present invention is implemented. The system comprises a source unit 101 and a receiving unit 102. Note
10 that the used terminology only reflects the roles the units 101-102 fulfil with respect to the present invention. The source unit 101 may e.g. be the mobile services switching centre MSC1 in Fig. 1 and the receiving unit 102 may be one of the base stations, e.g. BS1, connected to the mobile
15 services switching centre MSC1.

The receiving unit 102 comprises first storage means 104 for storing configuring data 105 and the source unit 101 comprises second storage means 106 for storing corresponding source configuring data 107. The configuring data 105 and
20 the source configuring data 107 are arranged in one or several groups of data 113. The source unit 101 and the receiving unit 102 are interconnected via data transfer means 103, i.e. some kind of communication network 108 and matching interfaces 109-110 at the respective unit 101-102,
25 enabling transfer of configuring data and other information between the source unit 101 and the receiving unit 102. The source unit 101 further comprises control means 111, e.g. a processor or a group of processors executing software, controlling the operation of the source unit 101. Likewise,
30 the receiving unit 102 also comprises control means 112, e.g. a processor or a group of processors executing software, for controlling the operation of the receiving unit 102.

In the distributed system 100 it is desirable that the configuring data 105 at the receiving unit 102 is consistent with the corresponding source configuring data 107 at the source unit 101. Due to disturbances of the receiving unit
5 102, the configuring data 105 may become corrupted and thus inconsistent with the source configuring data 107.

Fig. 2 is a flow chart illustrating a basic method according to the invention for synchronizing the configuring data 105 at the receiving unit 102 with the corresponding source
10 configuring data 107 at the source unit 101.

At step 201 reference checksums are calculated for each data group.

At step 202 the receiving unit 102 determines whether the content in each data group 113 of the configuring data 105
15 at the receiving unit 102 matches the corresponding reference checksum calculated at step 201.

At step 203 the receiving unit 102 requests the source unit 101 to transfer copies of the source configuring data 107 in those data groups 113 for which a mismatch with respect to
20 the corresponding reference checksum was found at step 202.

At step 204 the requested data, i.e. copies of the source configuring data in those data groups for which a mismatch was found at step 202, is downloaded from the source unit 101 to the receiving unit 102.

25 The basic method demonstrated in Fig. 2 illustrates some important characteristics of the inventive method.

One basic characteristic of the method is that when corruption of the configuring data 105 is detected, only copies of those data groups considered to be corrupted need
30 to be downloaded. This can amount to considerable savings in terms of the amount of data being downloaded from the source

1998 -12- 15

9

unit 101 to the receiving unit 102 if there is a great number of data groups of which e.g. only one has become corrupted.

Another basic characteristic of the method is that the receiving unit 102 determines whether any data group 113 of the configuring data 105 at the receiving unit 102 have become corrupted, i.e. the source unit 101 is relieved of this task. Thus the processing load at the source unit 101 is reduced which is an important advantage, especially if the source unit 101 is connected to a plurality of receiving units. As an example, a typical Ericsson CMS8800 mobile services switching centre may as a source unit be connected to 100-150 base stations.

There are a number of different alternative ways of implementing the basic method illustrated in Fig. 2 resulting in different embodiments of the inventive method.

The reference checksum calculations at step 201 may preferably be performed by the source unit 101 using the content of the source configuring data 107 at the source unit and downloaded to the receiving unit 102 prior to performing step 202. The reference checksums may be calculated according to the well known cyclic redundancy-code (CRC) algorithm. As a person skilled in the art realises, there are several other alternative ways of calculating reference checksums apart from using the CRC-algorithm. One example of an alternative way of calculating reference checksums would be to divide each data group into a number of data words, determine a parity bit for each data word in the data group so that the data word together with the parity bit contains an even (or alternatively odd) number of ones and combine all the parity bits determined for the data group into a reference checksum for the data group. Another example of an alternative way of calculating reference checksums would be to divide each data group into

a number of data words, determine N number of control bits for each data word as the value of the data word modulo 2^N and combine all the control bits determined for the data group into a reference checksum for the data group.

5 Instead of performing the reference checksum calculations using the content of the data groups 113 of the source configuring data 107, the reference checksum calculations at step 201 may be performed using the content of the data groups 113 of the configuring data 105 at the receiving unit
10 102 e.g. upon initial downloading of the configuring data 105. This latter approach implies that corruption of the configuring data 105 occurring after calculation of the reference checksums can be detected but not corruption occurring e.g. in connection with initial downloading of the
15 configuring data 105.

One way of performing step 202 is to divide step 202 into two substeps:

- a) calculating checksums at the receiving unit 102 for each data group 113 of the configuring data 105 stored in the
20 first storage means 105 at the receiving unit 102;
- b) comparing the calculated checksums to the corresponding reference checksums.

Steps 202-204, may be performed repeatedly as indicated by the dashed line in Fig. 2. This implies that the receiving
25 unit 102 repeatedly performs monitoring cycles to determine whether the configuring data 105 has become corrupted and requests updating of the configuring data 105 in the event of corruption. By repeatedly monitoring the configuring data consistency and performing redeeming actions in the event of
30 corrupted data being detected, the configuring data 105 at the receiving unit 106 is reliable to a higher extent and the need for downloading of new configuring data at system restart is reduced. Another significant advantage of the invention when repeatedly monitoring the configuring data

consistency, is that the source unit 101 does not need to calculate the reference checksums anew for each monitoring cycle, thus reducing the processing load of the source unit 101.

5 The monitoring cycles at step 202 may be initiated periodically with a predetermined time interval between each monitoring cycle. The frequency of initiating monitoring cycles at step 202 is selected as a tradeoff between the requirement of fast detection and correction of a data
10 inconsistency and the requirement to minimize the processing load caused when performing a monitoring cycle.

Instead of monitoring all data groups 113 of the configuring data 105 at the receiving unit in each monitoring cycles, the data groups 113 of the configuring data 105 at the
15 receiving unit 102 may be divided into a number of subsets and monitoring cycles may be initiated with individual periodicity for each subset. The predetermined time intervals used between each monitoring cycle may thus be selected differently for the different subsets. This
20 alternative may be attractive e.g. in a situation when data groups are distributed over a number of different memory units at the receiving unit of which some are more likely to experience disturbances than others.

The source unit 101 may initiate performance of step 202 by
25 sending an explicit signal to the receiving unit 102. However, when repeatedly performing monitoring cycles, the receiving unit 102 preferably initiates performance of step 202 on its own. Thus the processing load of the source unit 101 is further reduced. These two different ways of
30 initiating step 201 can also be combined. As an example of such a combination, the receiving unit 102 may periodically initiate monitoring cycles using reference checksums previously downloaded from the source unit 101. In addition to the monitoring cycles being periodically initiated by the

1998 -12- 15

12

receiving unit 102, the source unit 101 may initiate a complementary monitoring cycle by sending an explicit signal to the receiving unit 102 e.g. when the source unit 101 detects operation disturbances of the receiving unit 102.

- 5 The signal from the source unit 101 may contain fresh reference checksums to be used in the complementary monitoring cycle and subsequent periodically initiated monitoring cycles.

- 10 The different data groups 113 may be classified according to how urgently the content of each data group is needed by the receiving unit 102 in order to operate and step 204 may be performed so that copies of the source configuring data in data groups classified as more urgent are downloaded before copies of the source configuring data in data groups
- 15 classified as less urgent. In the context of the CMS8800 system, this feature could be applied so that e.g. data groups containing configuring data needed to enable operation of a control channel and at least one digital traffic channel or voice channel are classified as more
- 20 urgent and data groups containing data needed for supervision of the base station are classified as less urgent. Thus, the data needed to enable operation of the control channel and the at least one digital traffic channel or voice channel is downloaded before other data groups.
- 25 This enables the base station to restart operation and provide service earlier than if this urgent data was the last data downloaded.

- A preferred embodiment of the invention for use in the context of the cellular network NET1 in Fig. 3 is
- 30 illustrated in Fig. 4-7.

A schematic block diagram of one of the base stations BS1-BS5 in Fig. 3 is illustrated in Fig. 4. The base station BS1 comprises an interface 401 for connecting the base station BS1 via PCM-links to the mobile services swiching centre

MSC1. A number of extension modules 402 are connected in parallel to the interface 401. Each extension module 402 comprises an Extension Module Regional Process (EMRP) unit 403 and a number of devices 404 connected to the EMRP-unit 403. In the base station BS1 there are different kinds of devices such as Transceiver devices (TRX), Radio Frequency Test Loop Equipment (RFTL), Alarm module (ALM) and Combiner Tuner Controller (CTC). The devices 404 illustrated in Fig. 4 are transceivers, i.e. devices transmitting and receiving radio signals. The EMRP-unit 403 includes an Extension Module Regional processor 405 and an EMRP-memory unit 406. Each device 404 comprises a device processor 407 and a device memory unit 408.

The EMRP-memory units 406 in all extension modules together with all the device memory units 408 in the base station BS1 collectively correspond to the first storage means 104 in Fig. 1. In a similar way, all the Extension Module Regional processors 405 together with the device processors 407 collectively correspond to the control means 112 in Fig. 1.

Fig. 5A-5B illustrates how the configuring data at the base station BS1 in Fig. 4 as well as the corresponding source configuring data at the mobile services switching centre MSC1 are divided into a number of data groups. Note that Fig. 5A-5B do not disclose all the data groups in the base station BS1 but only an exemplary subset to illustrate the invention.

Fig. 5A illustrates two data groups which are stored in the EMRP-memory unit 406 in the base station BS1.

One data group 501 comprises cell configuring data such as information defining which cells are considered as neighbouring cells and parameters used in handoff algorithms for determining when a mobile station MS1-MS3 should be handed over to/from the cell C1 served by the base station BS1.

A second data group 502 comprises data used in the supervision of the extension module 402. This data group 502 comprises configuring data 505 defining the frequency of initiating monitoring cycles for determining the consistency of the data groups stored in the EMRP-memory unit 406 and the respective device memory unit 408 in the extension module 402. The data group 502 also comprises reference checksums 506-507 corresponding to the data groups 501-502 stored in said memory units 406 and 408.

Fig 5B illustrates two different types of data groups 503 and 504 which are stored in each device memory unit 408 of each transceiver (TRX) 404 in the extension module 402.

One data group 503 comprises equipment data such as:

- number of digital traffic channels supported by the TRX 404;
- channel number defining the radio carrier frequencies used when transmitting and receiving radio signals from/to the TRX;
- permitted maximal output power level used when transmitting radio signals;

The other type of data group 504 comprises information defining digital traffic channels. There is one instance of this data group for each digital traffic channel supported by the TRX. Each individual data group of this type comprises various data such as:

- initial power level to be used by a mobile station when communicating on the digital traffic channel;
- signal strength thresholds defining when a mobile station should be ordered to decrease or increase its transmission power;
- a signal strength threshold defining when communication with the mobile station should be blocked due to insufficient signal strength;
- Digital Verification Color Code;

One important aspect when defining the data groups is to consider how the configuring data will be distributed on different modules in the base station. Each data group should preferably be stored in a single memory unit, i.e. a device memory unit 408 or an EMRP-memory unit 406. The reason for this is primarily to provide for a more efficient way of calculating checksums on the content of the data groups.

Fig. 6 is a functional block diagram illustrating function blocks related to the present invention in the mobile services switching centre MSC1 and the base station BS1 in Fig. 3.

The mobile services switching centre MSC1, which acts as a source unit, comprises the function blocks coordinator-source 601, reference checksum calculator 602 and download controller-source 603. The base station BS1, which acts as a receiving unit, comprises the function blocks coordinator-receiving 604, checksum calculator 605 and download controller-receiving 606.

The coordinator-source function block 601 interacts with the other function blocks 602-603 in the mobile services switching centre MSC1 and is responsible for coordinating all activities related to the present invention within the mobile services switching centre MSC1. The coordinator-source function block 601 also interacts with the coordinator-receiving function block 604 at the base station BS1. The reference checksum calculator 602 calculates reference checksums for each data group of source configuring data stored at the mobile services switching centre MSC1 upon receiving an order from the controller-source 601 to perform such calculations. The download controller-source 603 is responsible for copying data from specified data groups of the source configuring data at the mobile services switching centre MSC1 and for transferring

the copied data to the base station BS1. The download controller-source 603 is ordered by the coordinator-source 601 to initiate downloading of specified data groups.

5 The coordinator-receiving function block 604 interacts with the other function blocks 605-606 in the base station BS1 and is responsible for coordinating all activities related to the present invention within the base station BS1. The coordinator-receiving function block 604 also interacts with the coordinator -source function block 601 at the mobile
10 services switching centre MSC1. The checksum calculator 605 calculates checksums for each data group of configuring data stored at the base station BS1 upon receiving an order from the coordinator-receiving function block 604 to perform said calculations. The download controller-receiving function
15 block 606 is responsible for receiving the data transferred by the download controller-source function block 603 and storing said data at the base station BS1.

In the base station BS1, the function block coordinator-receiving 604 is distributed on each EMRP-unit 403, the
20 function block checksum calculator 605 is distributed on each EMRP-unit 403 and each device 404, and the function block download controller-receiving 606 is essentially distributed on each EMRP-unit 403. As a consequence of each extension module 402 being essentially independent from the
25 other extension modules, the mobile services switching centre MSC1 interacts with each individual extension module 402. Thus when downloading data, the mobile services switching centre MSC1 downloads to each individual extension module 402 data being relevant for the respective extension
30 module 402, each extension module 402 individually notifying the mobile services switching centre MSC1 as corrupted data groups are detected.

Fig. 7 is a flow chart illustrating a preferred embodiment of a method according to the invention for continuously

keeping configuring data at the base station BS1 consistent with the corresponding source configuring data stored at the mobile services switching centre MSC1. The preferred embodiment of the inventive method will be described with reference to Fig. 7 and also to Fig. 4 and Fig. 6.

At step 701 the coordinator-source function block 601 in the mobile services switching centre MSC1 orders the reference checksum calculator 602 to calculate reference checksums for each data group of the source configuring data related to the base station BS1. At step 702 the coordinator-source function block 601 orders the download controller-source 603 to download all relevant source configuring data related to the base station BS1. The download controller-source 603 then retrieves all said relevant source configuring data, including the reference checksums, and transfers this data in update messages 607 to the base station BS1 via the PCM-links interconnecting the mobile services switching centre MSC1 and the base station BS1. One update message 607 is sent for each individual extension module 402 in the base station BS1. The download controller-receiving block 606 receives said data and stores the different data groups in the appropriate memory units at the base station BS1. Thus, after completion of step 702, data groups 501 and 502 are stored in the EMRP-memory unit 406 and the data groups 503-504 are stored in the device memory unit 408.

At the base station BS1, monitoring cycles are performed in parallel on a per extension module basis, i.e. the configuring data at the base station BS1 is divided into subsets corresponding to each extension module and monitoring cycles are performed in each extension module. This implies that steps 703-708 are performed in parallel for each extension module 402.

At step 703 the coordinator receiving function block 604 waits until it is time to initiate a monitoring cycle for the extension module 402. When it is time to initiate a monitoring cycle, the coordinator-receiving block 604 orders

the checksum calculator block 605 to perform checksum calculations for each data group of the configuring data stored within the extension module 402. At step 704, the checksums are calculated in parallel by the extension module

5 regional processor 405 and each device processor 407 in the extension module for the data groups stored in the respective memory unit 406, 408 connected to the respective processor 405, 407. The calculated checksums are reported to the coordinator-receiving block 604. At step 705 the

10 calculated checksums are compared by the coordinator-receiving block 604 to the corresponding reference checksums 506-507 of the second data group 502 stored in the EMRP memory unit 406. If there is no data group for which the calculated checksum is different from the corresponding

15 reference checksum (an alternative NO), processing continues at step 703. If a mismatch is found (an alternative YES) between the calculated checksum and the corresponding reference checksum for any of the data groups, e.g. the data group 503, processing continues at step 707 where the

20 coordinator-receiving block 604 issues a status report message 608 to the coordinator-source block 601 informing the coordinator-source block 601 that the data group 503 of the configuring data at the base station BS1 is considered to be corrupted, thus requesting the mobile services

25 switching centre MSC1 to download a fresh copy of the source configuring data in this data group to the base station BS1. At step 708 the coordinator-source block 601 orders the download controller-source block 603 to retrieve a copy of the source configuring data in this data group and to

30 transfer said copy in an update message 607 to the base station BS1. The download controller-receiving block 606 receives said data and updates the device memory unit 408 with the new copy of the data group 503. After step 708 processing continues at step 703, when the coordinator-receiving block 604 waits a predetermined time, i.e. in

35 accordance with the configuring data 505 defining the

monitoring frequency, before initiating a new monitoring cycle by ordering checksum calculations at step 704 etc.

In addition to the method illustrated in Fig. 7, a method enabling the mobile services switching centre MSC1 to
5 initiate individual monitoring cycles upon suspecting data in the base station to be corrupted may be provided. Thus, the coordinator-source function block 601 of Fig. 6 may send a synchronization signal 609 to the coordinator-receiving
10 block 604 after which, in essence, method steps 704-708 in Fig. 7 are performed. The synchronization signal 609 may contain fresh reference checksums to be used when performing the comparisons at step 705. Alternatively the reference checksums may be downloaded using an update message 607.

Apart from the exemplifying embodiments of the invention
15 disclosed above, there are several ways of providing rearrangements, modifications and substitutions resulting in additional embodiments of the invention.

As a person skilled in the art appreciates, the invention is in no way limited to CMS8800 cellular system. The invention
20 is applicable to a wide range of other distributed systems and in particular in distributed systems where a source unit is connected to a plurality of receiving units and/or synchronization is performed repeatedly. The source unit and/or the receiving unit may in themselves be distributed
25 systems.

A receiving unit may be connected to more than one source unit, thus containing data groups of configuring data whose corresponding source configuring data groups are stored in different source units. The receiving unit thus contains
30 different sets of configuring data where the corresponding source configuring data is stored in different source units. The receiving unit may be aware of what source unit the

different sets of configuring data originates from and interact directly with each respective source unit.

Alternatively, the receiving unit may be unaware that there are different sets of configuring data originating from
5 different source units. The receiving unit may thus treat the configuring data as if originating solely from a first source units of the source units, i.e. interacting only with the first source unit. The first source unit may in turn act as a relay between the receiving unit and the other source
10 units by forwarding data from the other source units to the receiving unit and forwarding requests from the receiving unit for fresh copies of source configuring data groups stored at the other source units.

1998 -12- 15

CLAIMS

1. A method for synchronizing configuring data (105) at a receiving unit (102) with corresponding source configuring data (107) stored at a source unit (101), the configuring data (105) and the source configuring data (107) each being arranged in at least one group of data (113), the method comprises the steps of:
 - a) calculating (201) reference checksums for each data group (113);
 - 10 b) determining (202) whether the content in each data group (113) of the configuring data (105) at the receiving unit (102) matches the corresponding reference checksum;
 - c) downloading (204) copies of the source configuring data (107) in those data groups (113) for which a mismatch was found at step b) from the source unit (101) to the receiving unit (102),
 - 15 c h a r a c t e r i z e d in that step b) is performed at the receiving unit (102) and that the method further comprises the step of:
 - 20 d) requesting (203) the source unit (101) to transfer copies to the receiving unit (102) of the source configuring data (107) in those data groups (113) for which a mismatch is found at step b), wherein step d) is performed between step b) and step c).
- 25 2. A method according to claim 1, wherein the reference checksums are calculated using the content of the source configuring data (107) at the source unit (101) and the method further comprises the step of:
 - e) downloading the calculated reference checksums to the receiving unit (102), wherein step e) is performed prior to
 - 30 step b).
3. A method according to any one of claims 1-2, wherein step b) is initiated upon detecting operation disturbances of the receiving unit (102).

4. A method according to any one of claims 1-2, wherein step b) is performed repeatedly.

5. A method according to claim 4, wherein step b) is initiated periodically with a predetermined time interval
5 between each cycle (202).

6. A method according to claim 4, wherein the data groups (113) of the configuring data (105) at the receiving unit (102) are divided into at least two subsets and step b) is initiated periodically for each subset with predetermined
10 time intervals between each cycle (202), the predetermined time intervals being selected individually for each respective subset.

7. A method according to any one of claims 1-6, wherein step b) comprises the steps of:
15 f) performing checksum calculations (704) for each data group (113) of the configuring data (105) at the receiving unit (102);
g) comparing (705) the calculated checksums to the corresponding reference checksums.

20 8. A method according to any one of claims 1-7, wherein the data groups (501-504) are classified according to the urgency of the content of each data group with respect to the operation of the receiving unit (102) and step c) is performed so that copies of the source configuring data
25 (107) in data groups (503, 504, 501) classified as more urgent are downloaded prior to downloading copies of the source configuring data (107) in data groups (502) classified as less urgent.

9. A distributed system (100) comprising a receiving unit
30 (102), a source unit (101) and data transfer means (103) interconnecting the receiving unit (102) and the source unit (101), wherein the receiving unit (102) includes first

storage means (104) for storing configuring data (105) and the source unit (101) includes second storage means (106) for storing corresponding source configuring data (107), the configuring data (105) and the source configuring data (107) each being arranged in at least one group of data (113), the distributed system (100) comprises:

reference checksum calculating means (602) for calculating reference checksums for each data group (113);

determining means (604, 605) for determining whether the content in each data group (113) of the configuring data (105) at the receiving unit (102) matches the corresponding reference checksum;

downloading means (603, 606) for downloading copies from the source unit (101) to the receiving unit (102) of the source configuring data (107) in those data groups (113) for which the determining means (604, 605) has found a mismatch between the content of the configuring data (105) at the receiving unit (102) and the corresponding reference checksums,

characterized in that the determining means (604, 605) is located at the receiving unit (102) and that the distributed system (100) comprises means (604) for requesting the source unit (101) to download copies of the source configuring data (107) in those data groups (113) for which the determining means (604, 605) found a mismatch.

10. A distributed system (100) according to claim 9, wherein the reference checksum calculating means (602) is located in the source unit (101) and is adapted to calculate the reference checksums using the content of the source configuring data (107) stored in the second storage means (106), and the downloading means (603, 606) is adapted to download the calculated reference checksums from the source unit (101) to the receiving unit (102).

11. A distributed system (100) according to any one of claims 9-10, wherein the determining means (604, 605) is adapted to determine whether the content in each data group (113) of the configuring data (105) at the receiving unit (102) matches the corresponding reference checksum (113) upon detection of operation disturbances of the receiving unit (102).

12. A distributed system (100) according to any one of claims 9-10, wherein the determining means (604, 605) is adapted to repeatedly perform monitoring cycles (202) determining whether the content in each data group (113) of the configuring data (105) at the receiving unit (102) matches the corresponding reference checksum (113).

13. A distributed system (100) according to claim 12 wherein the determining means (604, 605) is adapted to periodically initiate the monitoring cycles (202) with a predetermined time interval between each monitoring cycle (202).

14. A distributed system (100) according to claim 12 wherein the data groups (113) of the configuring data (105) at the receiving unit (102) are divided into at least two subsets and the determining means (604, 605) is adapted to periodically initiate the monitoring cycles (202) for each subset with predetermined time intervals between each monitoring cycle (202), the predetermined time intervals being selected individually for each respective subset.

15. A distributed system (100) according to any one of claims 9-14 wherein the determining means (604, 605) comprises:
checksum calculating means (605) for calculating checksums for each data group (113) of the configuring data (105) at the receiving unit (102);

comparing means (604) for comparing the checksums calculated by the checksum calculating means (605) to the corresponding reference checksums.

5 16. A distributed system (100) according to any one of claims 9-15 wherein the data groups (501, 502, 503, 504) are classified according to the urgency of the content of each data group with respect to the operation of the receiving unit (102) and the downloading means (603, 606) is adopted
10 to download copies of the source configuring data (107) in data groups (503, 504, 501) classified as more urgent prior to downloading copies of the source configuring data (107) in data groups (502) classified as less urgent.

ABSTRACT

A method and arrangement in a distributed system for synchronizing configuring data at a receiving unit with corresponding source configuring data. The configuring data and the source configuring data are arranged in one or preferably several groups of data. Reference checksums are calculated (201) for each data group. The receiving unit determines (202) whether the content in each data group of the configuring data at the receiving unit matches the corresponding reference checksum. The receiving unit requests (203) the source unit to transfer copies to the receiving unit of the source configuring data in those data groups for which a mismatch is detected. The requested data is downloaded (204) from the source unit to the receiving unit.

Fig. 2 for publication.

1998 -12- 15

1/7

100

102




Fig. 1

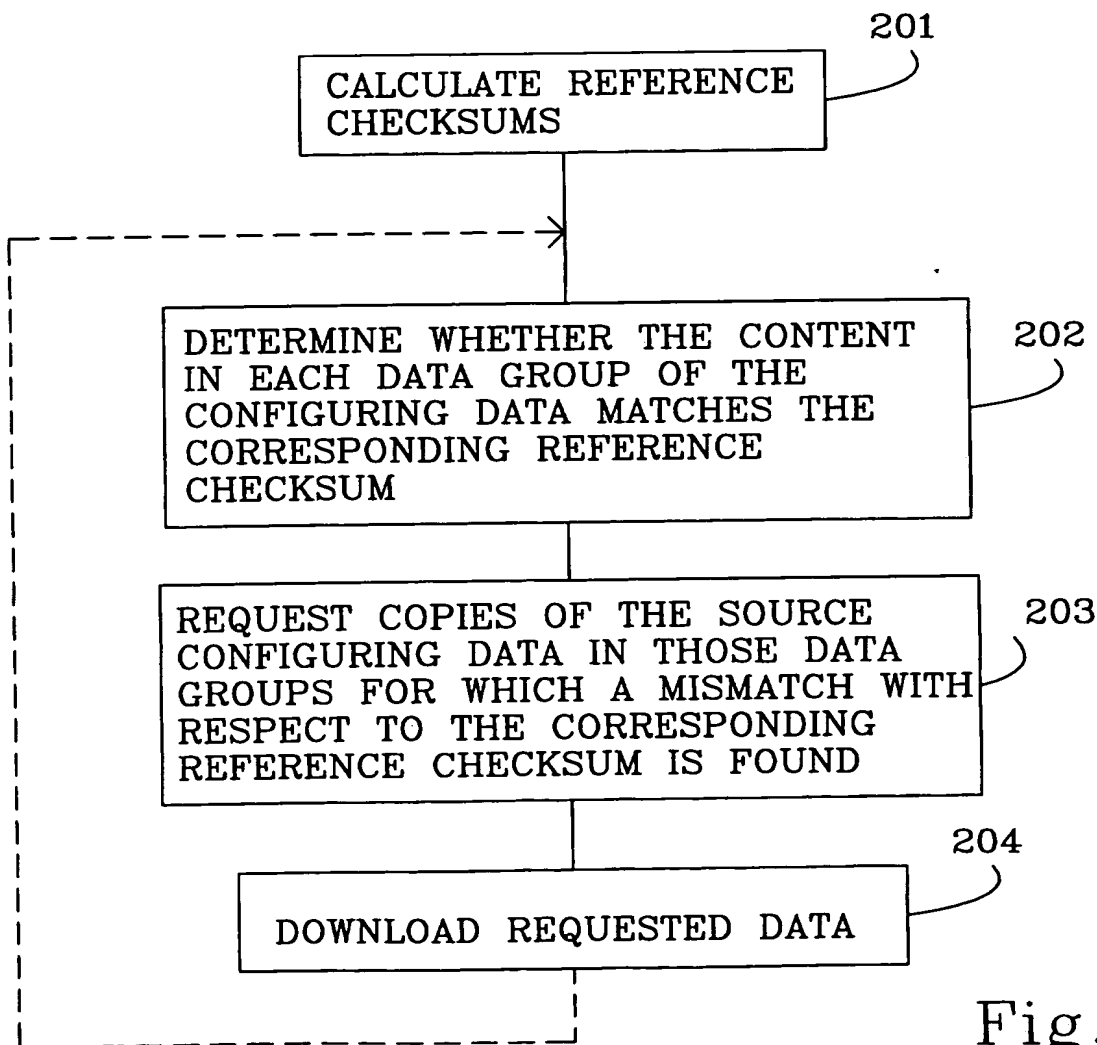


Fig. 2

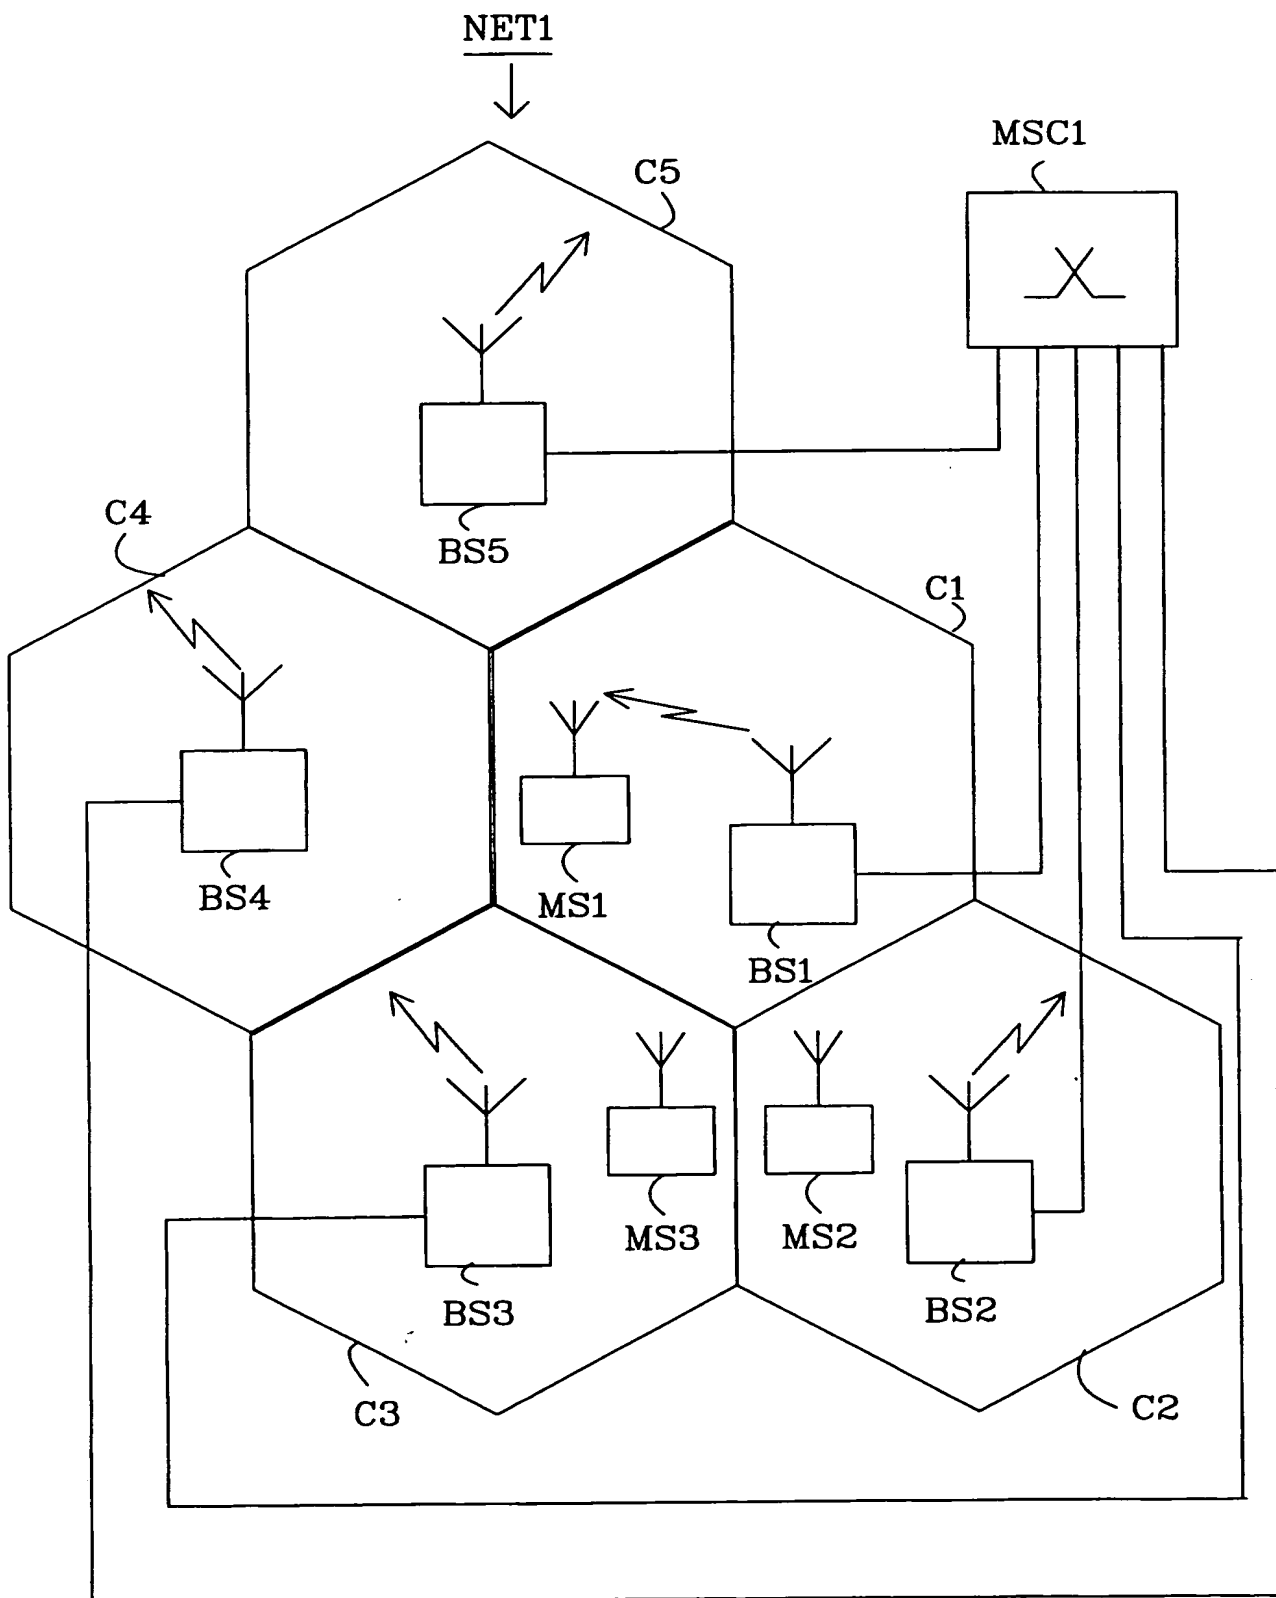


Fig. 3

3/7

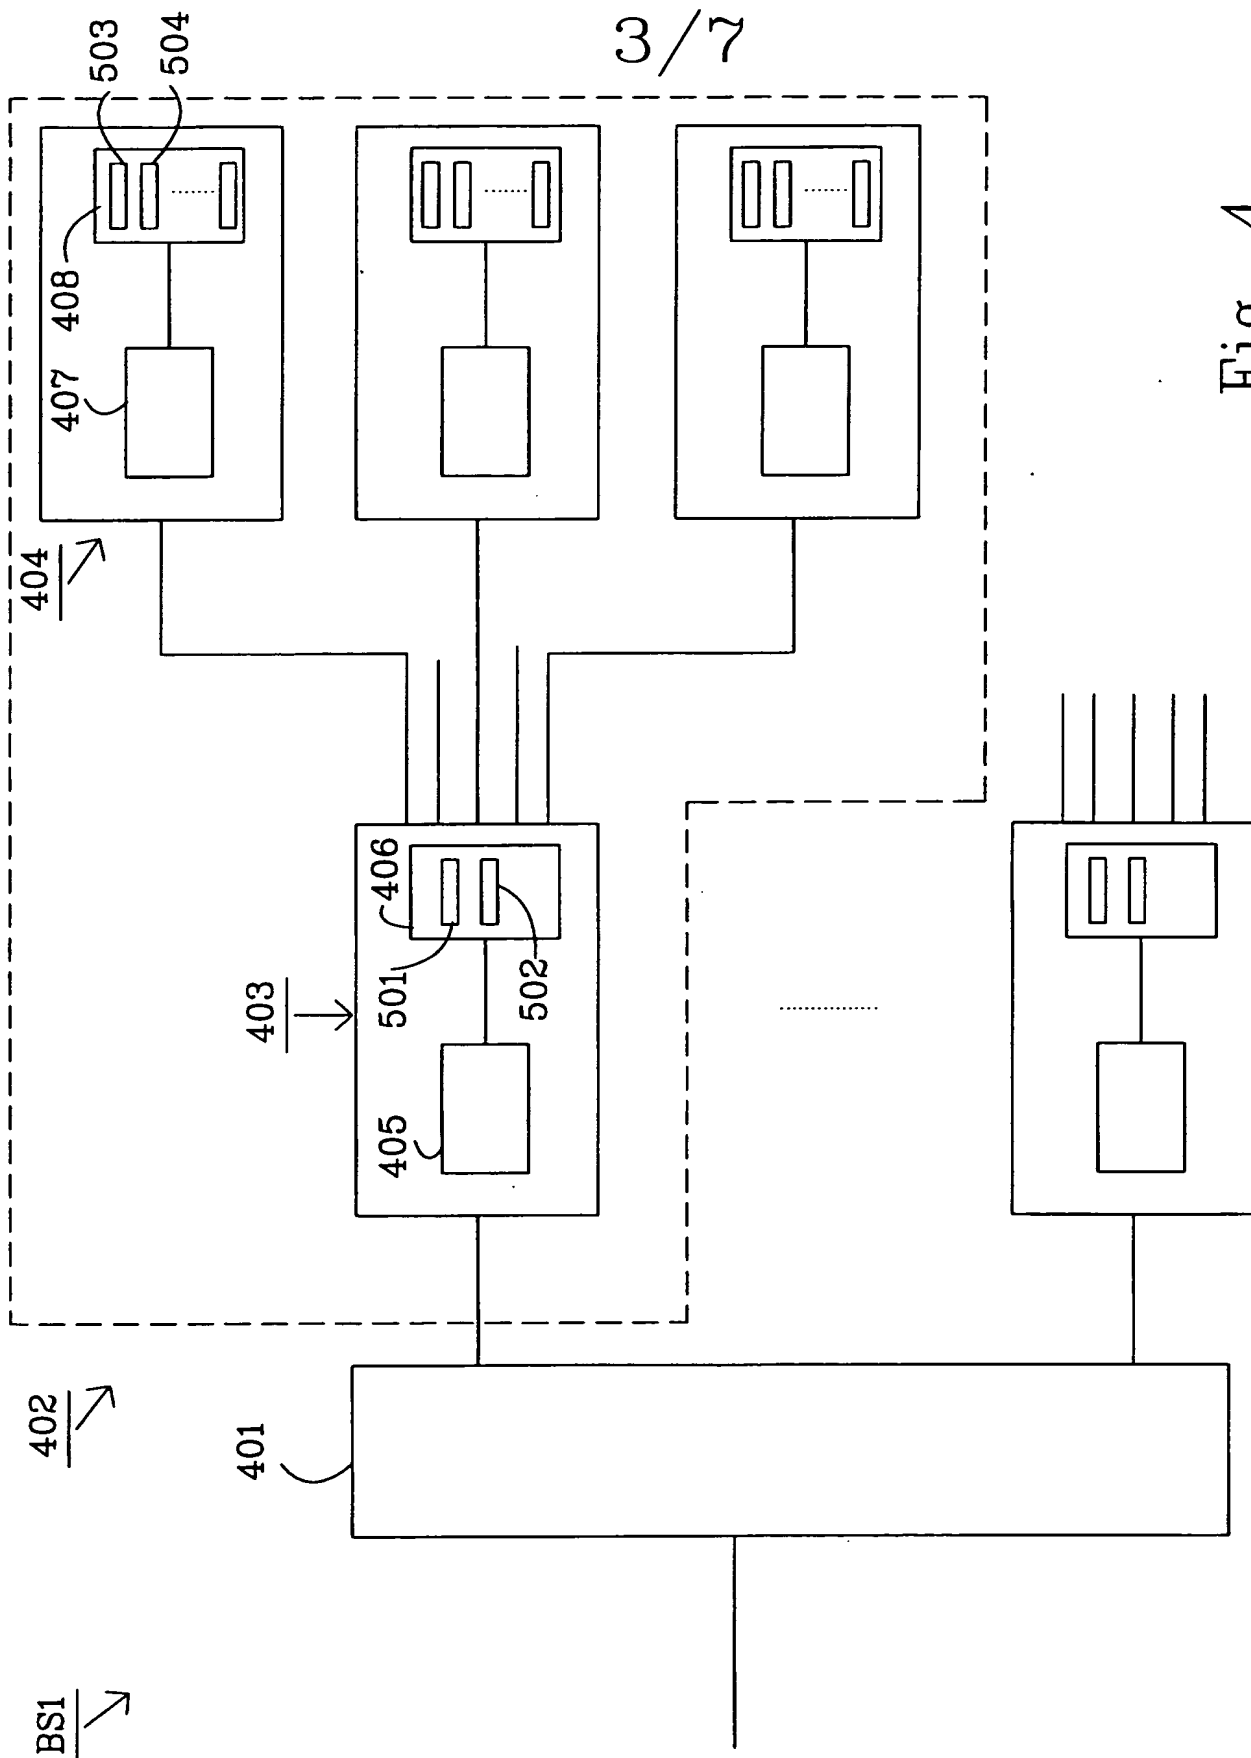


Fig. 4

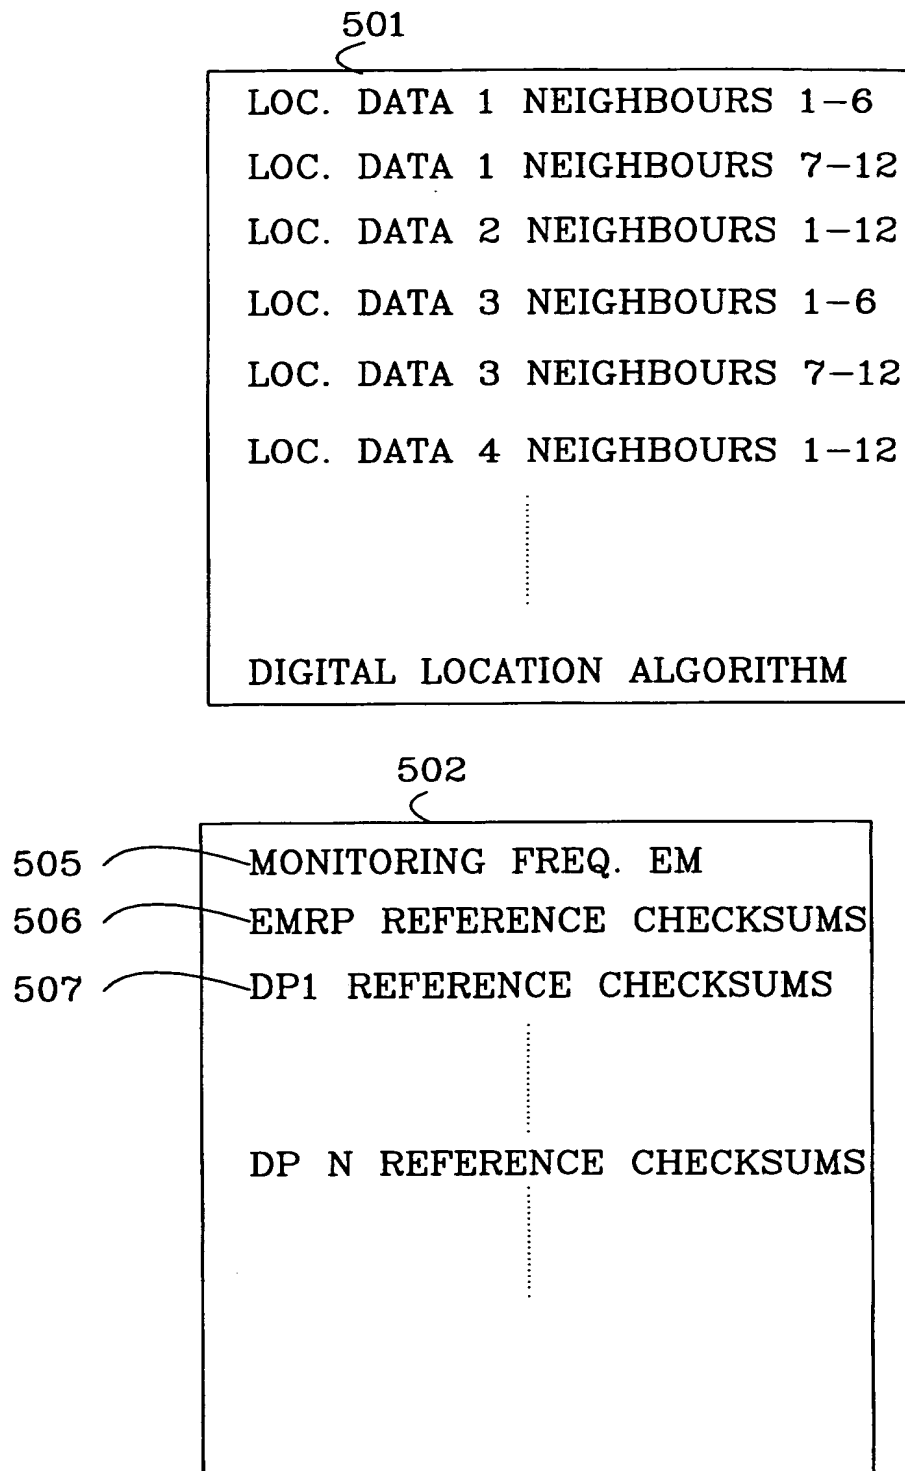


Fig. 5A

1998 -12- 15

5/7




Fig. 5B

1998 -12- 15

6/7

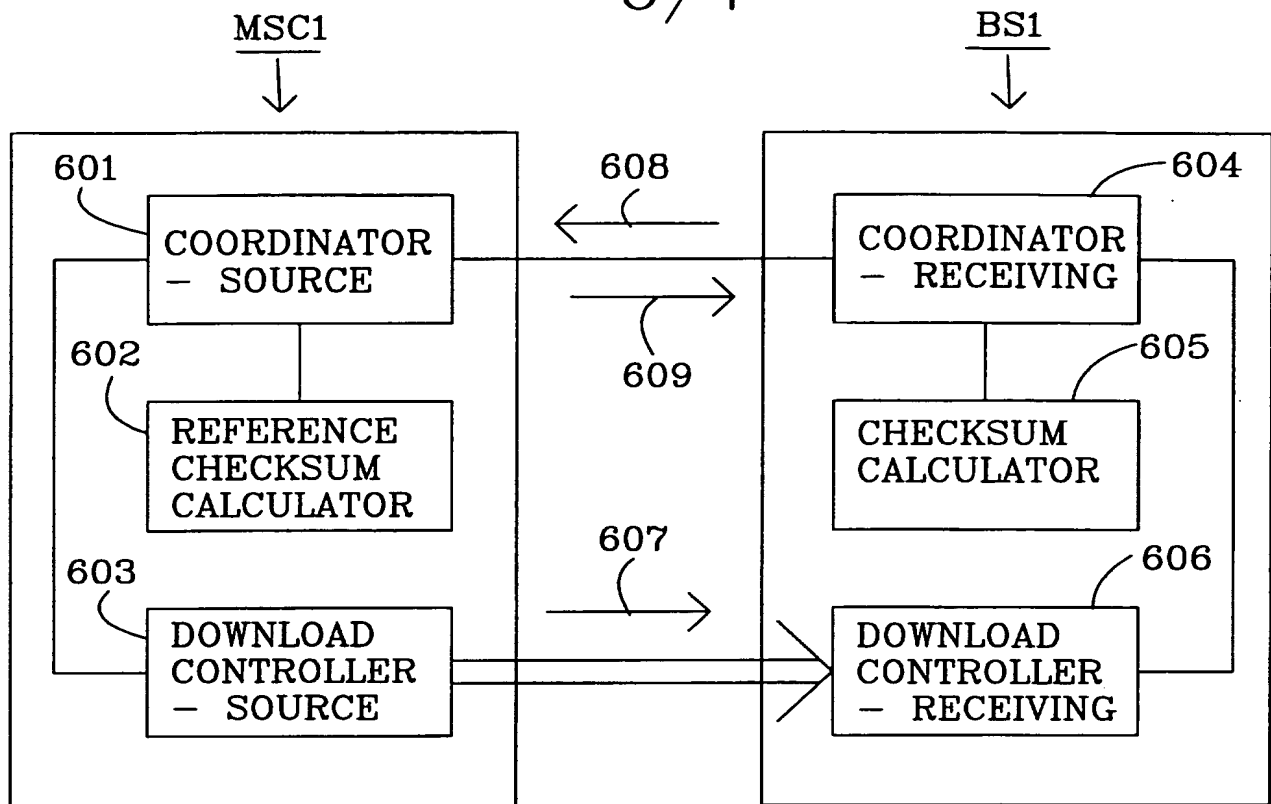


Fig. 6

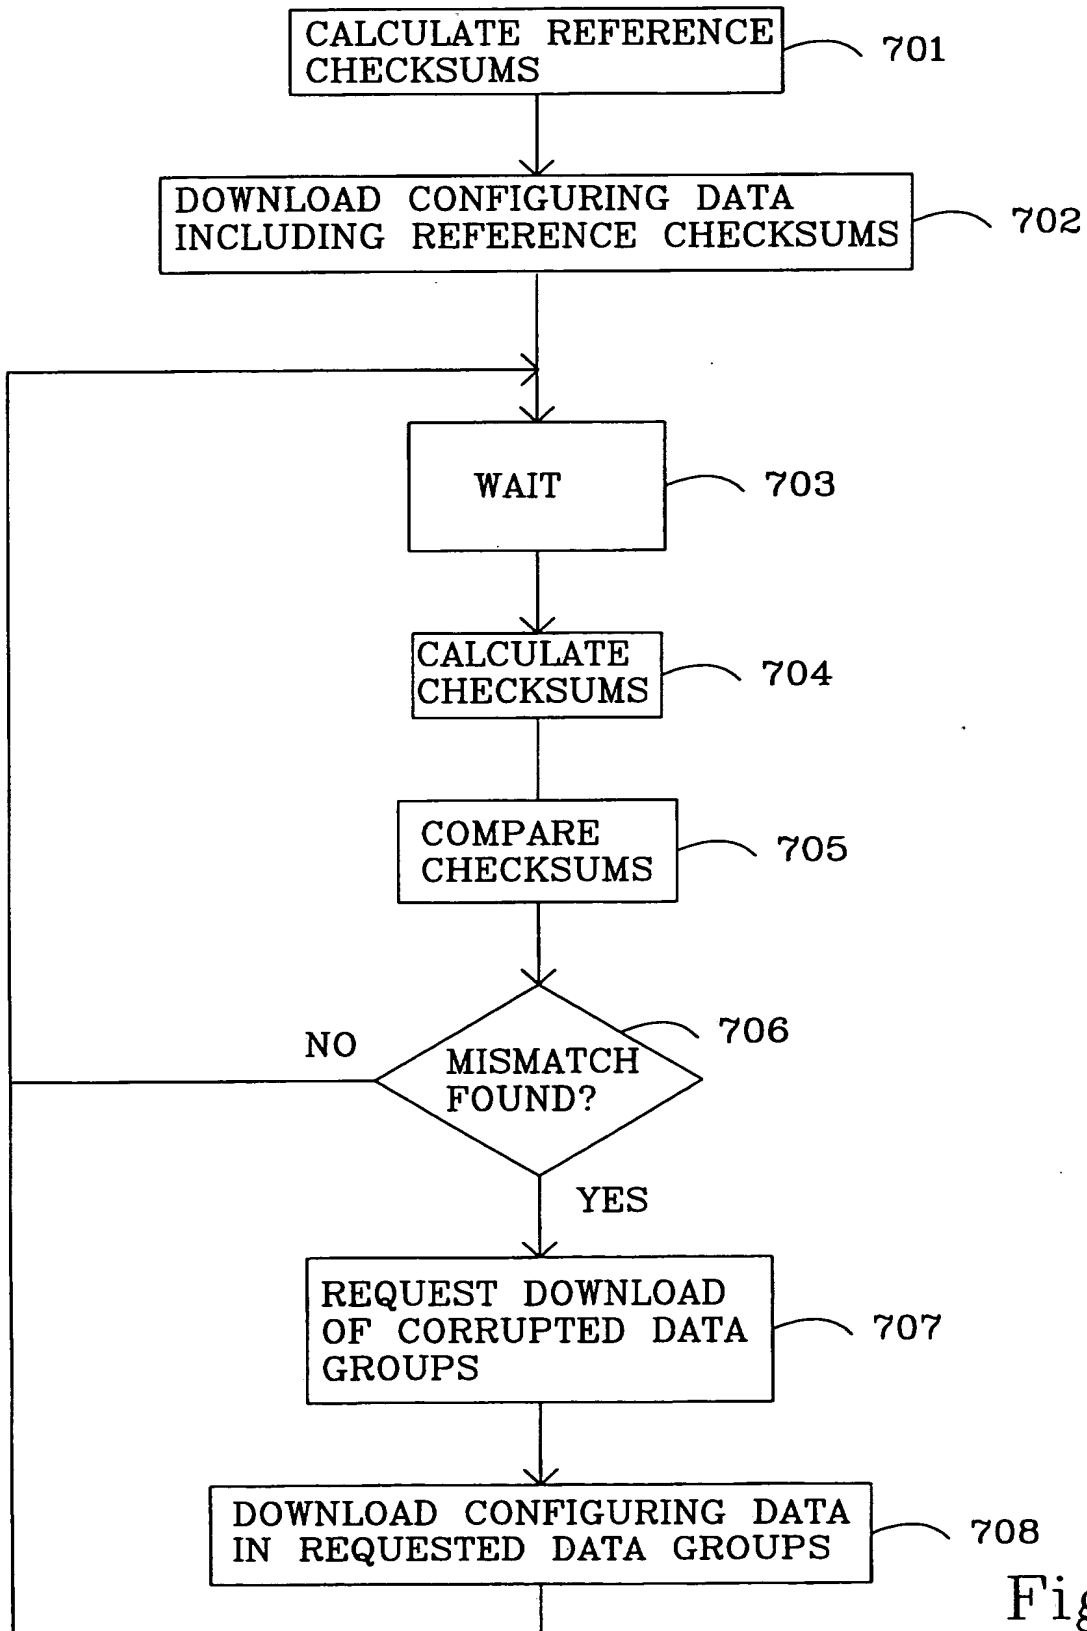


Fig. 7